Modeling and Analysis of Call Center Arrival Data:
A Bayesian Approach

Refik Soyer[*]
Department of Management Science
The George Washington University

M. Murat Tarimcilar
Department of Management Science
The George Washington University

In this paper we present a modulated Poisson process model to describe and analyze arrival data to a call center. The attractive feature of this model is that it takes into account both covariate and time effects on the call volume intensity and in so doing enables us to assess the effectiveness of different advertising strategies along with predicting the arrival patterns. A Bayesian analysis of the model is developed and an extension of the model is presented to describe potential heterogeneity in arrival patterns. The proposed model and the methodology are implemented using real call center arrival data.

Key words: call center; advertising strategy; modulated Poisson process; Bayesian analysis; heterogeneity; random effects

---

1. Introduction and Background

Forecasting call volume data that is based on a specific advertisement or promotion plan poses a challenging problem for call centers. The need for optimal scheduling and staffing of telephone operators in call centers has made forecasting models an important component of decision making in many organizations [see for example, Andrews and Parsons (1993) and Gans et al. (2003)]. Some of the earlier forecasting

[*]Contact author: Department of Management Science, Monroe Hall 403, The George Washington University, Washington, DC 20052, USA. E-mail: soyer@gwu.edu

models include ARIMA processes and transfer function models. These have been successfully used in forecasting call volumes [see for example Andrews and Cunningham (1995)]. More recent approaches involve use of queuing models as in Jongbloed and Koole (2001) and doubly stochastic Poisson models as in Avramidis et al. (2004). The focus of these previous work was to model the call arrival (demand) process based on aggregate arrival data.

In addition to modeling call center arrival process, it is desirable to evaluate the efficiency of and return on advertisement and promotion policies to develop marketing strategies for call centers. This requires an advertisement specific analysis of arrival patterns rather than analysis of aggregate arrival data. In a recent project we were faced with a need to develop models and methods for assessing the impact of individual advertising campaigns of print media on call arrivals. The predictions produced by such models have significant implications for marketing and advertising budget decisions as well as for scheduling and staffing. The models and the statistical methodology presented here are motivated by this project.

Modeling and analysis of call center arrival data from a marketing point of view need to address the following issues:

(a) development of a model describing the call volume over time for different advertising and promotion policies by accounting for covariate effects as well as the life cycle of advertisements;

(b) assessment of the effectiveness of different advertising strategies by accounting for the effect of other relevant covariates;

(c) ability to describe heterogeneity in advertisements beyond what is captured by covariates;

(d) ability to predict call volume over any time interval so that optimal scheduling of operators can be determined.

2

In this paper we present a modulated Poisson process model that provides us with a framework to address the above issues and develop its Bayesian analysis. Unlike the previous models used in call arrival analysis, the proposed model allows for advertisement specific analysis of call arrival data and thus enables us to assess effectiveness of various advertisement and promotion strategies. In modeling daily call arrival data Avramidis et. al. (2004) have noted that arrivals during different time partitions are correlated. They have proposed Poisson models with stochastic arrival intensities to deal with this problem. In our development, by taking a Bayesian approach we describe our uncertainty about parameters of the modulated Poisson process model probabilistically via specifying a prior distribution. This approach yields a model that can be considered as a doubly stochastic Poisson model as used in Avramidis et. al. (2004) and this results in correlated call arrival counts. Thus, the Bayesian treatment of the modulated Poisson process model provides an alternate modeling approach to describe correlated call arrival data. Furthermore, this introduces a methodology for statistical analysis of call center operations as Gans et al. (2003) alluded to as a potential research prospect.

In recent years the Bayesian approach have been considered in many marketing related problems; see Rossi and Allenby (2003) for a review. As pointed out by the authors, the Bayesian methods provide a more flexible framework in dealing with heterogeneity. As discussed in Allenby and Rossi (1999), Bayesian models have been used in dealing with consumer heterogeneity but they have not been not considered in the analysis of call center arrival data. In this paper, as an extension of the modulated Poisson process model, we introduce a random effects type model to describe potential heterogeneity in advertisements and present Bayesian inference for the model.

In what follows, we present the call arrival data and introduce a modulated Poisson process model to describe the arrival pattern of calls over time by accounting for covariate effects. This is done in Section 2. This model enables us to assess the effectiveness of different advertising venue, of various advertising strategies and provides

a framework for predicting call volume over time. In Section 3 we present a Bayesian analysis of the model and discuss how posterior and predictive distributions are obtained in such an analysis. To consider heterogeneity in different advertisements, a random effects type extension of the model is considered in Section 3.1 and a model comparison approach is discussed Section 3.2. An illustrative example is given in Section 4 using real call volume center data. Conclusions are presented in Section 5.

2. A Model for Call Center Arrivals

2.1 Call Arrival Data

The data we used in our analysis comes from a consumer electronics producer who offers a limited variety of products. Significant portion of the sales goes through the call center. Products usually have long life cycles and the life cycle is usually extended by periodic updates and upgrades. With the aging product line, the advertisement budget has been increasing drastically.

The data is more detailed than what is commonly available from a typical call center. The time arrival history is available for each advertisement separately along with advertisement specific data such as cost, frequency and type. The company issues advertising in most media venues however majority is in print media. Each advertisement is targeted and aims at urging the customer to place the call. The company tracks the origin of the call to the specific advertisement that led to the call, becaues each incoming call is stamped with advertisement ID. The advertisement ID typically points out the cost of advertisement, its medium, the advertisement format, and the type of promotion being offered. The incoming number of calls for each advertisement is available as interval counts for different periods over the life cycle of the advertisement. An example of call arrival data for a typical advertisement is shown in Table 1 below.

4

Table 1: An Example of Call Arrival Data for an Advertisement

Time Interval (in days)	Number of Calls
(0, 1]	6
(1, 2]	5
(2, 3]	1
(3, 4]	3
(4, 5]	2
(5, 6]	2
(6, 7]	2
(7, 8]	0
(8, 9]	2
(9, 10]	2
(10, 18]	0

2.2 Modulated Poisson Process Model

In modeling call arrivals generated by an advertisiment venue such as in Table 1, it is important to note that the effectiveness of the advertisement decreases by time. In view of this fact, the number of calls generated by an advertisement in a specified time interval can be described by a nonhomogeneous Poisson process (NHPP). As pointed out earlier, it is of interest to assess the effectiveness of different advertising strategies as well as advertising venues on call volume. Thus, these factors as well as the time effect should be considered by any model describing the call volume. Modeling the effect of covariates on the call volume intensity requires a modulation of the Poisson process as proposed in Cox (1972a).

Let $N_i(t)$ denote the number of calls arrive during a time interval of length t as response to the ith advertisement and let Z_i denote a $p \times 1$ vector of covariates that describe the characteristics of the ith advertisement. Typically, Z_i will consist of components such as media expense (in \$'s), venue type (monthly magazine, daily newspaper etc.), advertisement format (full page, half page, color etc.), offer type (free shipment, payment schedule etc.) and seasonal indicators.

To reflect the fact that effectiveness of advertisement i is a function of time, $N_i(t)$ is described by a NHPP with intensity function

$$\lambda_i(t) = \frac{d}{dt} E[N_i(t)], \qquad (2.1)$$

where $E[\bullet]$ denotes the expectation. The modulated Poisson process model (MPPM) assumes that the intensity function of the ith advertisement is related to the covariate vector $\boldsymbol{Z}_i$ via

$$\lambda_i(t, \boldsymbol{Z}_i) = \lambda_0(t) e^{\boldsymbol{\beta}' \boldsymbol{Z}_i}, \qquad (2.2)$$

where $\lambda_0(t)$ is the baseline intensity function and $\boldsymbol{\beta}$ is a $p \times 1$ vector of parameters. Note that (2.2) incorporates both the time and covariate effects on call intensity.

The MPPM's have been considered in survival analysis by Sinha (1993). The modulated Poisson process model can be thought as a counting process alternative to the proportional hazards model (PHM) of Cox (1972b) and is sometimes referred to as proportional intensities model (PIM) in reliability modeling literature [Merrick et. al. (2004)]. Similar to the PHM, in the PIM ratio of intensity functions

$$\frac{\lambda_i(t, \boldsymbol{Z}_i)}{\lambda_j(t, \boldsymbol{Z}_j)} = e^{\boldsymbol{\beta}'(\boldsymbol{Z}_i - \boldsymbol{Z}_j)} \qquad (2.3)$$

does not depend on time. The cumulative intensity (or the mean-value) function of the MPPM is given by

$$\Lambda_i(t, \boldsymbol{Z}_i) = \Lambda_0(t) e^{\boldsymbol{\beta}' \boldsymbol{Z}_i}, \qquad (2.4)$$

where $\Lambda_0(t) = \int_0^t \lambda_0(s) ds$. Given $\Lambda_0(t)$, $\boldsymbol{\beta}$, and $\boldsymbol{Z}_i$, the distribution of $N_i(t)$, call volume in any interval of length t, is given by

$$P(N_i(t) = n | \Lambda_0(t), \boldsymbol{\beta}, \boldsymbol{Z}_i) = \frac{(\Lambda_0(t) e^{\boldsymbol{\beta}' \boldsymbol{Z}_i})^n}{n!} exp\left(-\Lambda_0(t) e^{\boldsymbol{\beta}' \boldsymbol{Z}_i}\right). \qquad (2.5)$$

Similarly, conditional on $\Lambda_0(\bullet)$, $\boldsymbol{\beta}$, and $\boldsymbol{Z}_i$, the probability distribution of distribution of number of calls in a time interval $(s, t]$ is given by the Poisson model

$$P(N_i(t) - N_i(s) = n | \Lambda_0(\bullet), \boldsymbol{\beta}, \boldsymbol{Z}_i) = \frac{\left(\Lambda_i(t, \boldsymbol{Z}_i) - \Lambda_i(s, \boldsymbol{Z}_i)\right)^n}{n!}$$

$$\times \, exp\left\{ - [\Lambda_i(t, \boldsymbol{Z}_i) - \Lambda_i(s, \boldsymbol{Z}_i)] \right\}. \tag{2.6}$$

In modeling the baseline intensity $\Lambda_0(t)$ of the modulated Poisson Process, a parametric form can be specified. For example, an appropriate form for $\Lambda_0(t)$ in the analysis of call arrival data can be the *power law model* which is also used in reliability and survival analysis. This is given by

$$\Lambda_0(t) = \gamma \, t^\alpha \tag{2.7}$$

where $\alpha > 0$ and $\gamma > 0$. In the power law model, values of $\alpha < 1$ imply that the effectiveness of an advertisement deteriorates with time. This is typically what is expected in call volume generated by a given advertisement. The specification of the above form also implies that the distribution of the time to the first call arrival is a Weibull density with shape parameter α. For this reason the Poisson process with the power law intensity function is referred to as *Weibull process.*

We can investigate the appropriateness of the power law model by obtaining a scatter plot of cumulative number of calls against time in the log scale. Taking logs on both sides of (2.7) yields

$$log[\Lambda_0(t)] = log\gamma + \alpha \, logt. \tag{2.8}$$

The above implies that if the plot of log of the expected number of calls against the log of time is approximately linear, then the power law model is appropriate. In Figure 1, we show such a scatter plot for a group of monthly advertisements that are in similar cost

7

range. The linear fit that is shown in the figure suggests that the parametric form (2.7) is reasonable for the cumulative intensity function. In our illustration in Section 4, we will use the power law as the intensity function of the modulated Poisson process model. However, it is possible to consider other parametric forms for $\Lambda_0(\bullet)$ or to use a nonparametric setup as discussed in Gelfand (1999).
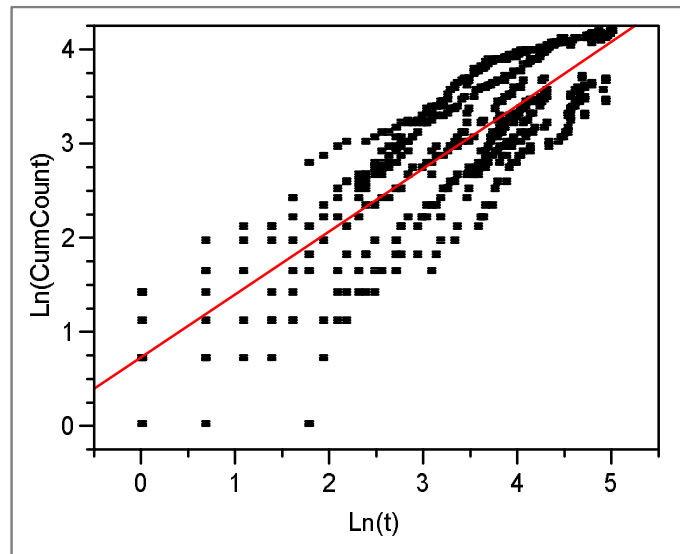

Figure 1. Scatter plot of Cumulative Number of Calls versus Time in Log Scale.

Recent studies in modeling aggregate daily arrivals to call centers have noted the inappropriateness of the standard Poisson process models with deterministic arrival intensities. Avramidis et. al. (2004) point out that arrival counts in nonoverlapping time partitions are correlated and propose alternate stochastic models for arrival counts.

In what follows, we will present a Bayesian analysis of the MPPM whose baseline intensity function is given by the power law model (2.7). The Bayesian approach requires that uncertainty about all unknown parameters of the MPPM, that is, α, γ and covariate parameter vector $\boldsymbol{\beta}$, is described probabilistically by specifying a prior distribution $p(\alpha, \gamma, \boldsymbol{\beta})$. Note that the results (2.5) and (2.6), as well as the independent increments

property of the NHPP, hold only conditional on the parameters α, γ and β. In other words, unconditionally the $N_i(t)$ process will not have independent increments. Thus, the Bayesian treatment of the MPPM provides an alternate modeling framework to alleviate difficulties arise in analysis of call arrival data.

3. Bayesian Analysis of the Modulated Poisson Process Model

We assume that the call volume data is available in terms of interval arrivals as shown in Table 1, but the presented approach can easily be modified to incorporate time of arrivals data. For advertisement i, the process $N_i(t)$ is observed at r_i time intervals with endpoints

$$t = t_{i1}, \ldots, t_{ir_i}, \text{ and } t_{i1} < \ldots < t_{ir_i},$$

$i = 1, \ldots, m$ and the call volume data generated by advertisement i is given by

$$D_i = \{N_i(t_{ij}) = n_i(t_{ij}), \, j = 1, \ldots, r_i, \boldsymbol{Z}_i\},$$

where $\boldsymbol{Z}_i$ denotes the covariate vector associated with advertisement i. Using the independent increments property of the NHPP conditional on the parameters, given data on advertisement i, the likelihood function of $\Lambda_0(\bullet)$ and $\boldsymbol{\beta}$ is given by

$$L_i(\Lambda_0(\bullet), \boldsymbol{\beta}; D_i) = \prod_{j=1}^{r_i} \frac{\left([\Lambda_0(t_{ij}) - \Lambda_0(t_{i,j-1})]e^{\boldsymbol{\beta}'\boldsymbol{Z}_i}\right)^{n_i(t_{ij}) - n(t_{i,j-1})}}{[n_i(t_{ij}) - n(t_{i,j-1})]!}$$

$$\times exp\left\{-\left([\Lambda_0(t_{ij}) - \Lambda_0(t_{i,j-1})]e^{\boldsymbol{\beta}'\boldsymbol{Z}_i}\right)\right\}. \tag{3.1}$$

Conditional on $\Lambda_0(\bullet)$ and $\boldsymbol{\beta}$, the $N_i(t)$'s are assumed to be independent. Thus, given the call counts for each $N_i(t)$ at $t_{i1}, \ldots, t_{ir_i}$, where $t_{i1} < \ldots < t_{ir_i}$, for $i = 1, \ldots, m$, the likelihood function of $\Lambda_0(\bullet)$ and $\boldsymbol{\beta}$ is

$$L(\Lambda_0(\bullet), \boldsymbol{\beta}; D) = \prod_{i=1}^{m} L_i(\Lambda_0(\bullet), \boldsymbol{\beta}; D_i) \tag{3.2}$$

where $D = \bigcup_{i=1}^{m} D_i$.

We will assume the power law model for the intensity function $\Lambda_0(\bullet)$. The Bayesian analysis requires specification of a joint prior distribution $p(\alpha, \gamma, \beta)$ for the power (α), and scale (γ) parameters of the intensity function and for β, the parameters of the covariate vector β. Note that the values and signs of the coefficients in the β vector provide information about the effectiveness of a covariate in generating calls. For example, if the covariate is the media expense, then a positive value of the coefficient implies that higher levels of advertising expenses result in a higher level of call intensity. In our development we will assume independent priors on these parameters. More specifically uncertainty about the covariate vector β will be described by a multivariate normal prior, denoted as $\beta \sim N(\mu, \Sigma)$, with a specified mean vector μ and a covariance matrix Σ. In specifying the prior parameters of this distribution we can incorporate any expert judgment on the effect of media expenses, advertising strategy, etc. into the analysis; see Lindley (1983) for a formal framework for use of expert opinion in Bayesian analysis.

In specifying the prior distribution of the parameters in the power law model, we note the fact that both α and γ are nonnegative random variables and furthermore values of $\alpha < 1$ (> 1) reflect the fact that the effectiveness of the advertisements deteriorate (improve) by time. In our development we will specify a gamma prior on α and a lognormal prior on γ. Again the prior parameters of these distributions can be chosen in such a way to reflect our knowledge or ignorance about them.

Given the above choice of priors or any other form, a fully Bayesian analysis of the MPPM can not be obtained analytically. In other words, the joint posterior distribution

$$p(\alpha, \gamma, \beta | D) \propto L(\Lambda_0(\bullet), \beta; D) \, p(\alpha, \gamma, \beta) \qquad (3.3)$$

is not in an analytically tractable form. However, a Gibbs sampler similar to that used in Dellaportes and Smith (1993) for analysis of PHM can be developed to simulate from the posterior distribution (3.3). Once we have samples from the posterior distribution we can

obtain predictions of type (2.5) and (2.6). The Gibbs sampler enables us to simulate from (3.3) by recursively sampling from full conditional distributions $p(\alpha|\gamma, \boldsymbol{\beta}, D)$, $p(\gamma| \alpha, \boldsymbol{\beta}, D)$ and $p(\boldsymbol{\beta}|\alpha, \gamma, D)$; see Gelfand and Smith (1990) for a review of the Gibbs sampler and implementation details. Given the above selection of priors, the full conditional distributions for α, γ and $\boldsymbol{\beta}$ can not be obtained in any known form. Thus, sampling from these requires use of special simulation algorithms. However, it can be easily shown that all of these distributions are logconcave. Thus, the adaptive rejection algorithm of Gilks and Wild (1992) can be used for sampling from these distributions.

Given the posterior sample

$$\left\{\alpha^{(l)}, \gamma^{(l)}, \boldsymbol{\beta}^{(l)}\right\}_{l=1}^{S}$$

from the joint posterior we can make call arrival predictions for any advertisement for any time interval. Note that the posterior predictive distribution of number of arrivals in an interval of t time units for advertisement i is given by

$$P(N_i(t) = n|D) = \int P(N_i(t) = n|\alpha, \gamma, \boldsymbol{\beta}, \boldsymbol{Z}_i) \, dp(\alpha, \gamma, \boldsymbol{\beta}|D), \qquad (3.4)$$

where $P(N_i(t) = n|\alpha, \gamma, \boldsymbol{\beta}, \boldsymbol{Z}_i)$ is given by (2.5). The above integral can not be evaluated analytically, but we can compute it using a Monte Carlo integral approximation as

$$P(N_i(t) = n|D) \simeq \frac{1}{S} \sum_{l=1}^{S} P(N_i(t) = n|\alpha^{(l)}, \gamma^{(l)}, \boldsymbol{\beta}^{(l)}, \boldsymbol{Z}_i) \qquad (3.5)$$

for $n = 0, 1, 2, \ldots$ Also, we can obtain expected number of calls in a time interval as

$$E(N_i(t)|D) \simeq \frac{1}{S} \sum_{l=1}^{S} E(N_i(t)|\alpha^{(l)}, \gamma^{(l)}, \boldsymbol{\beta}^{(l)}, \boldsymbol{Z}_i)$$

where

$$\mathrm{E}(N_i(t)|\alpha, \gamma, \boldsymbol{\beta}, \boldsymbol{Z}_i) = \Lambda_i(t, \boldsymbol{Z}_i) = \gamma t^{\alpha} e^{\boldsymbol{\beta}' \boldsymbol{Z}_i}.$$

11

Similarly, we can approximate the probability of n calls in the interval $(s, t]$, $s < t$, for advertisement i as

$$P(N_i(t) - N_i(s) = n|D) \simeq \frac{1}{S} \sum_{l=1}^{S} P(N_i(t) - N_i(s) = n|\alpha^{(l)}, \gamma^{(l)}, \boldsymbol{\beta}^{(l)}, \boldsymbol{Z}_i), \quad (3.6)$$

where $n = 0, 1, 2, \ldots$, and $P(N_i(t) - N_i(s) = n|\alpha, \gamma, \boldsymbol{\beta}, \boldsymbol{Z}_i)$ is given by (2.6). We can also evaluate the joint distribution of number of calls for any given type of advertisements using the conditional independence of $N_i(t)$'s given α, γ, $\boldsymbol{\beta}, \boldsymbol{Z}_i$, for $i = 1, 2, \ldots, m$.

3.1 Modeling Heterogeneity in Advertisements

The MPPM of (2.4) implies that differences between the intensity functions of advertisements are adequately captured by the included covariates of the model. It is possible that in many cases certain covariate information may not be available or the covariates included in the model will not reflect potential heterogeneity in different advertisements. Modeling heterogeneity in advertisement is essential not only for improving the predictive performance of the call arrival model but also to assess the effectiveness of different advertising and promotion strategies in creating the call volume.

Recently, Bayesian models have been used in marketing literature for modeling consumer heterogeneity; see for example, Allenby and Rossi (1999). However, these models have not included point process models such as the MPPM and they have not considered analysis of call center arrival data.

A common strategy to model heterogeneity is to consider a random effects type extension of the MPPM. This involves reparameterization of the intensity function as

$$\Lambda_0^i(t) = \gamma_i \, t^\alpha \qquad (3.7)$$

where the parameter γ_i, for $i = 1, 2, \ldots, m$, is assumed to be drawn from a known mixing distribution G. This type of extension results in a hierarchical Bayes type setup in

the sense of Lindley and Smith (1972). More specifically, in the intensity function (3.7), we define

$$log\gamma_i = \theta + \phi_i, \tag{3.8}$$

where ϕ_i's are the random effect terms. We assume that ϕ_i's are conditionally independent normal random variables

$$\phi_i|\tau \sim N(0, 1/\tau)$$

where the unknown precision is described by the gamma prior, $\tau \sim G(a_\tau, b_\tau)$, with specified parameters a_τ and b_τ. As before α will have a gamma prior and the coefficient vector $\boldsymbol{\beta}$ is assumed to follow a multivariate normal prior. Furthermore, in this setup we can also specify a normal prior for θ in (3.8), which can be incorporated into the multivariate normal prior of $\boldsymbol{\beta}$. In our setup, we assume that apriori $\boldsymbol{\beta}$, θ and α are independent of ϕ_i's and τ.

Under the random effect type model, the posterior distribution of interest is

$$p(\alpha, \theta, \boldsymbol{\beta}, \boldsymbol{\phi}, \tau|D) \propto \prod_{i=1}^{m} L_i(\Lambda_0^i(\bullet), \boldsymbol{\beta}; D_i)\, p(\alpha, \theta, \boldsymbol{\beta}, \boldsymbol{\phi}, \tau) \tag{3.9}$$

where $\boldsymbol{\phi} = (\phi_1, \ldots, \phi_m)$. As in the original model the posterior distribution (3.9) can not be obtained analytically. Thus, all the posterior and predictive analyses need to be done using a Gibbs sampler. In this case the required full conditionals also include $p(\tau|\boldsymbol{\phi}, \alpha, \boldsymbol{\beta}, D)$ and $p(\phi_i|\phi^{(-i)}, \tau, \alpha, \boldsymbol{\beta}, D)$, where $\phi^{(-i)} = \{\phi_j| j \neq i, j = 1, \ldots, n\}$, as well as the full conditionals of α and $\boldsymbol{\beta}$.

It follows from the above setup that the full conditional distribution of τ can be obtained as a gamma distribution given by

$$(\tau|\boldsymbol{\phi}, \alpha, \boldsymbol{\beta}, D) \sim G(a_\tau + n/2, b_\tau + \sum_{i=1}^{n} \phi_i^2/2).$$

The remaining full conditional distributions do not follow any standard forms. However, since they are all logconcave, the adaptive rejection sampling method can be used to

sample from these at each iteration of the Gibbs sampler. Similar to the original model, the predictive distributions (3.5) and (3.6) are obtained once posterior draws are available from the Gibbs sampler.

3.2 Model Comparison

In comparing the original modulated Poisson process model with its random effects extension presented above, computation of the Bayes factors [see Kass and Raftery (1995) for a comprehensive review] is difficult. The marginal likelihoods under the two competing models can not be directly approximated from the Gibbs sampler to compute the Bayes factor.

An alternative is to use a model selection criterion such as the *Deviance Information Criterion* (DIC) of Spiegelhalter et. al. (2002). For a generic parameter vector Θ, DIC is defined as

$$DIC = \overline{D} + p_D, \qquad (3.10)$$

where $D = -2log\mathcal{L}(\Theta)$, is two times the negative loglikelihood, $\overline{D} = E_{\Theta|data}[D]$ and $p_D = \overline{D} - D(\widehat{\Theta})$, where $\widehat{\Theta}$ is the posterior mean. The DIC has the general "fit + complexity" form used by many model selection criteria. In (3.10) $\overline{D}$ represents the "goodness of the fit of the model where p_D represents a complexity penalty as reflected by the effective number of parameters of the model. Note that p_D also includes the contribution of the random effects parameters to the model.

4. A Real Life Data Illustration

In what follows, we present an illustration of the model using the real call center data. The data used in the illustration is only a subset of the real data set we have received from the company. Due to company's specific and structured advertising policy, each call is traced to the advertisement that led the call. All the characteristics of the advertisement

are known, such as its cost, specials offered in the advertisement and its format. Our previous analysis indicated that the cost, which is named as media dollars, has a significant effect on the calls, both in its volume and its distribution. This variable is used as a covariate in the model. The results that are presented in this section are based on 142 different print media advertisements. Each advertisement has between 7 to 74 time intervals where call volumes are available. In our analysis we will use media dollars (in $000), that is, the cost for the advertisement, print media type (weekly or monthly) and the offer type as covariates. There will be a dummy variable which will capture the effect of media type and the monthly ads will be used as the reference group. There are three offer types and these are captured by two dummy variables in the model. The standard offer type will be used as the reference in the model and the first and second offer types represent interest bearing installment option and free originating and return shipment, respectively. In our model we also include an interaction term of cost and media type to allow for differences in cost effect of different media type.

We first consider the original modulated Poisson process model without the random effect terms and refer to this model as the fixed effects model in our discussion. We note that for the case where we include the offer type and cost of the advertising as covariates, log of the cumulative intensity function (2.4) can be written as

$$log[\Lambda_i(t, \boldsymbol{Z}_i)] = \theta + \alpha log t + \beta_1 Z_{1i} + \beta_2 Z_{2i} + \beta_3 Z_{1i} Z_{2i} + \beta_4 Z_{3i} + \beta_5 Z_{4i} \quad (4.1)$$

where $\theta = log\gamma$ in (2.7), Z_{1i} is cost of the advertisement, Z_{2i} is the dummy variable for media type and Z_{3i} and Z_{4i} are the dummy variables representing the second and third order types. The model also includes an interaction term of cost and media type. In (4.1) $\boldsymbol{\beta} = (\beta_1\ \beta_2\ \beta_3\ \beta_4\ \beta_5)$ is the covariate vector and for the monthly advertisements with standard offer type the cumulative intensity becomes

$$\Lambda_i(t, \boldsymbol{Z}_i) = \gamma t^\alpha exp(\beta_1 Z_{1i}). \qquad\qquad (4.2)$$

In our analysis flat but proper priors are used for all the unknown parameters and all the computations will be done using the BUGS programming environment of Spiegelhalter et al. (1996). In implementation of the Gibbs sampler, the results are obtained after an initial burn-in sample of 10,000 iterations. We have experienced no convergence problems and runs with different starting points gave us very similar results.

Unlike the classical approach, all inferences in Bayesian paradigm are described by posterior probability distributions of the unknown quantities. Summary statistics such as mean, mode and variance of the posterior distributions can also be reported. In Figure 2, we present the posterior distributions of α and β_1 of the fixed effects model. The posterior distribution of α which is concentrated in the interval (0.48, 0.52) indicating that the effectiveness of the advertisements deteriorates by time. In other words, for a fixed level of media expense and using a particular offer, as the advertisement ages, the expected number of calls during any time interval diminishes exponentially. This is an expected result and consistent with the advertisement literature. In Figure 2, The value of β_1 shows the impact of each $1,000 spent on the total number of calls received due to that specific advertisement. The positive values that the distribution concentrated around show that media expense implies an increase in the expected number of calls during any time interval for a given offer type.

The posterior distributions of β_2 and β_3 are shown in Figure 3. Both of these distributions are concentrated in the positive region, implying that weekly ads are more effective than the monthly ads in generating the call volume. Note that the positive values of β_3 represents a higher impact of ad cost for weekly ads on the expected number of calls received.
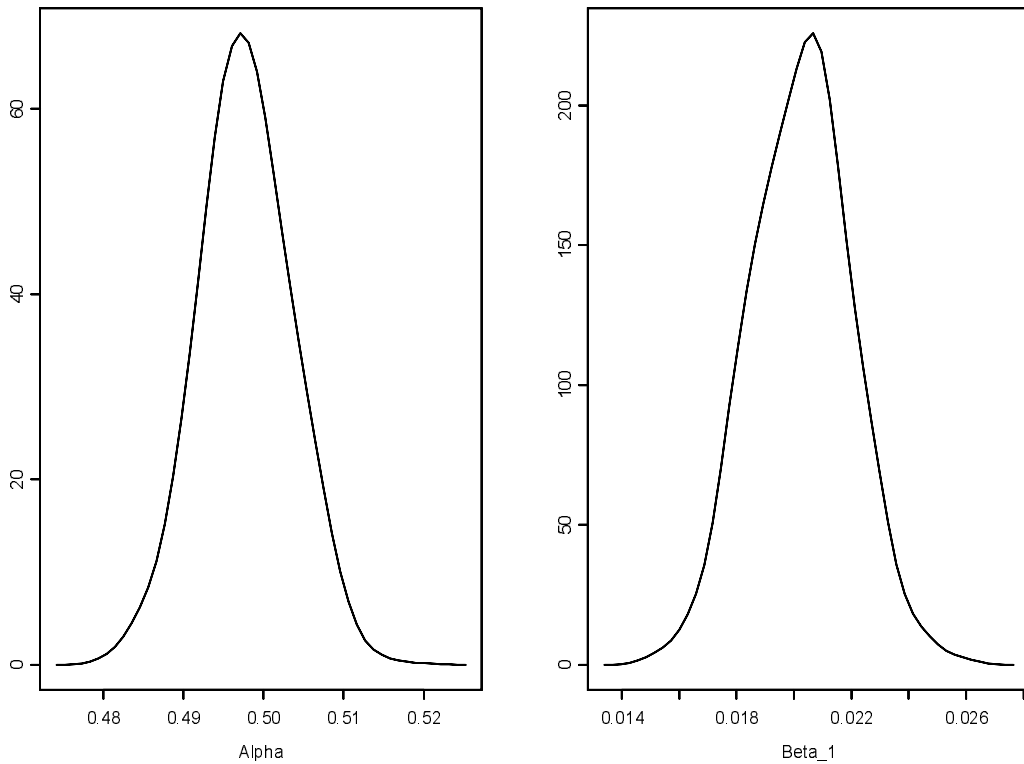
Figure 2. Posterior Distributions of α and β_1.

Posterior distributions of β_4 and β_5 that are shown in Figure 3 represent uncertainty about expected change in the call volume (in log scale) as a result of using a special offer. Based on the concentration of the posterior distribution of β_4 somewhere around $(0.1, 0.3)$, the expected change in call volume due to the interest bearing installment option seems to be positive relative to the standard offer. Similarly, as implied by the posterior distribution of β_5, the "free originating and return shipment" option also yields an expected increase in call volume beyond the standard offer.

Figure 4 displays the posterior predictive distributions of number of calls for four different time intervals for a monthly advertisement that costs \$4,500 and using the third offer type, that is, free originating and return shipment option. Specifically we present the distribution of number of calls during the periods of $0 - 5$, $5 - 10$, $10 - 15$ and $15 - 20$

17

days. As expected, the distribution of the number of calls becomes more concentrated around smaller values as we move from $0 - 5$ to $15 - 20$ days interval whereas the precision increases. We note that the posterior predictive distribution is helpful for the analyst particularly in setting the service level in call centers. Scheduling and staffing around the mean would lead to too many missed calls. By using the predictive distribution, the call center could run a resource allocation model keeping in mind the trade-offs between the opportunity cost of missed calls and the cost of staffing.
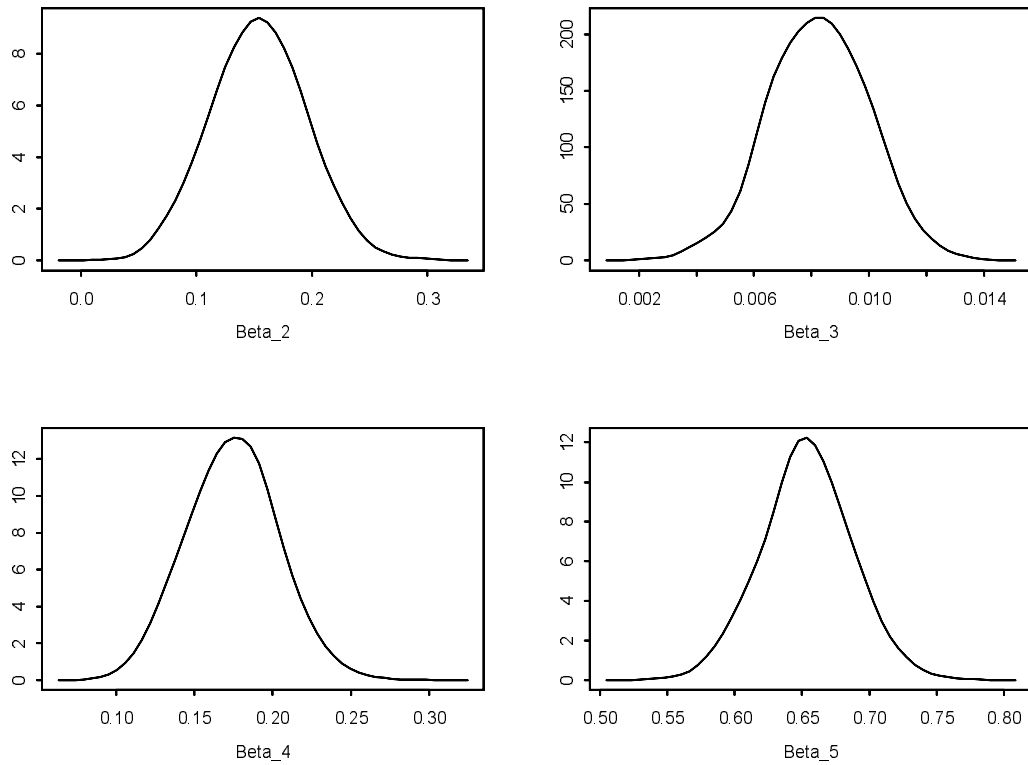

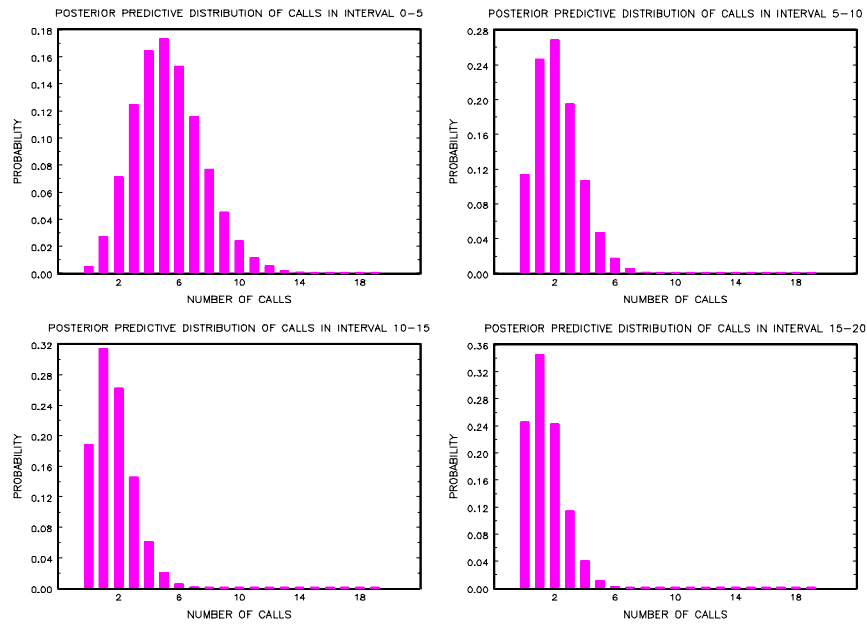Figure 3. Posterior Distributions of β_i's, $i = 2, \ldots, 5$.

Figure 4. Posterior Predictive Distributions of Number of Calls for a Monthly Ad with $Z_1 = \$4,500$ and Standard Offer Type under the Fixed Effect Model.

We have done a similar analysis using the random effects type model of Section 3.1. In so doing, we again used proper but noninformative priors in all cases including the random effect terms ϕ_i's and the precision parameter τ. As before an initial burn-in run of 10000 iterations have been used in the Gibbs sampler and 2000 samples have been obtained using a lag of 25 between successive samples. No convergence problems have been experienced in the Gibbs sampler. A comparison of the posterior means and standard deviations of the common parameters of the fixed and random effect models are given in Table 2. We note that in all cases, except the distribution of α, as expected, we have lot higher variability under the random effects model. The distribution of α is very similar in both cases. When we compare the distributions of β_5 under the two models, we see that the positive effect of the offer type two on call volume is more pronounced under the random effects model. On the other hand, the effect of the first offer type on expected call

19

volume, as implied by the distribution of β_4, is less pronounced. The posterior probability $Pr(\beta_4 > 0|D)$ can be obtained as 0.875 under the random effects model.

Table 2: Comparison of Posterior Means and Standard Deviations

	Fixed Effects Model		Random Effects Model	
	Mean	Std	Mean	Std
α	0.4977	0.0058	0.5012	0.0060
θ	0.7730	0.0484	0.5138	0.1593
β_1	0.0203	0.0017	0.0274	0.0069
β_2	0.1542	0.0408	0.0951	0.1601
β_3	0.0083	0.0017	0.0066	0.0082
β_4	0.1746	0.0289	0.1388	0.1231
β_5	0.6536	0.0331	0.8338	0.1459

Comparison of the posterior distributions of β_2 and β_3 show that the relative effectiveness of weekly ads in generating the call volume is less pronounced under the random effects model. The posterior probability $Pr(\beta_2 > 0|D)$ can be obtained as 0.726 under the random effects model whereas under the fixed effects model the posterior distribution was concentrated in the positive region. Similarly, under the random effects model, the posterior probability $Pr(\beta_3 > 0|D) = 0.795$ implies weaker evidence for the higher impact of ad cost for weekly ads on the expected number of calls received.

In Figure 5 we present the posterior distributions of random effect parameters ϕ_i's for thirty arbitrarily selected advertisements. We note that there are clear differences between these distributions. This suggests presence of heterogeneity in intensity of call arrivals generated by specific advertisements.

These differences also yield a more diffused posterior predictive distribution for the number of calls during the intervals of $0 - 5$, $5 - 10$, $10 - 15$ and $15 - 20$ days as shown in Figure 6. If we compare the posterior predictive distributions given in Figure 6 with those of Figure 4, we see that the posterior predictive distributions under the random effects model are more right skewed.

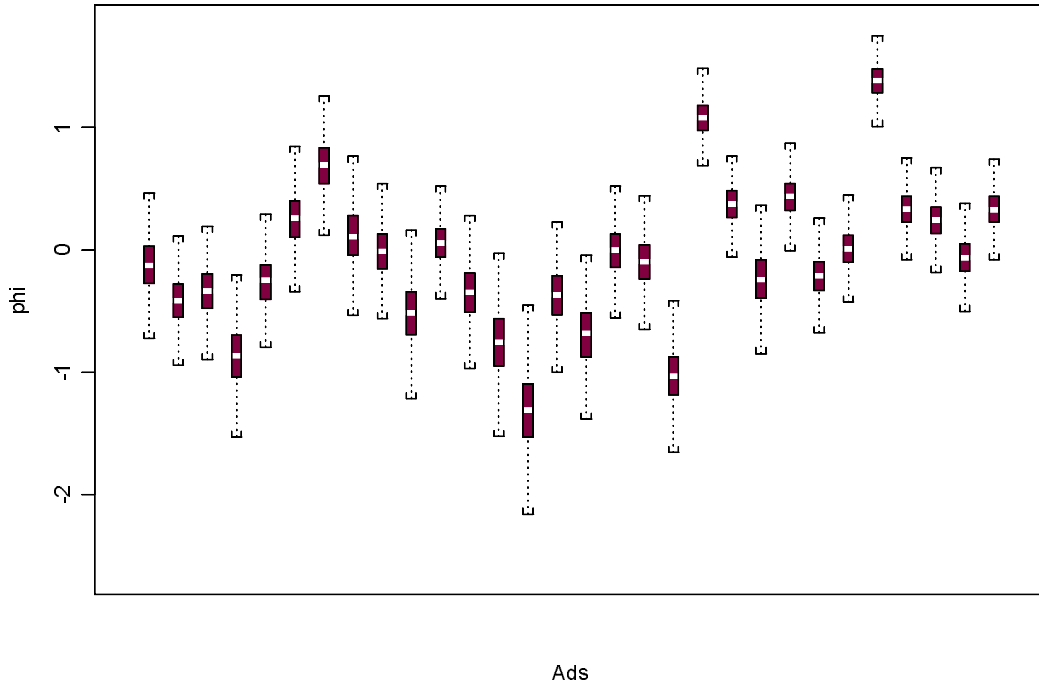Posterior Distribution of Random Effects Terms for Selected Ads

Figure 5. Posterior Distributions of ϕ_i's for Selected Advertisements.

A comparison of the two models is made using the DIC in Table 3. We note that the DIC under fixed effects model is lot higher than the one obtained under the random effects model. In other words, the data provides a strong support in favor of the random effects model. When we look at the estimated values of the effective number of parameters, that is, the p_D values, we see that in the fixed effects model this is 6.9 where the actual number is 7. On the other hand, under the random effects model the effective number of parameters is 130. We note that 8 of these are the actual parameters and the remaining ones represent the random effects. In other words, among the advertisements which we have analyzed, there are differences that can not be captured by the fixed covariates.
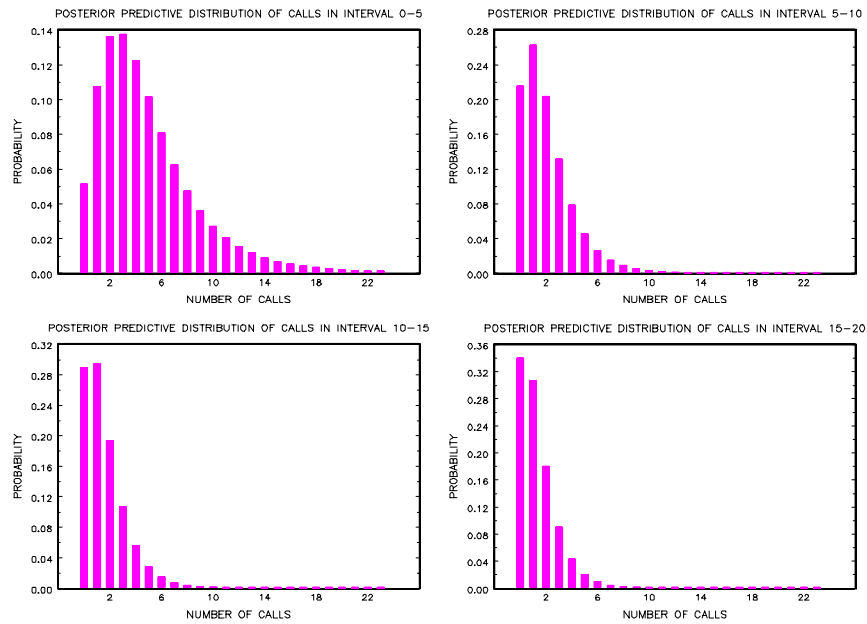
Figure 6. Posterior Predictive Distributions of Number of Calls for a Monthly Ad with $Z_1 = \$4{,}500$ and Standard Offer Type under the Random Effects Model.

Table 3: Comparison of the Models using DIC

Model	DIC	p_{D}
Fixed Effects	25430.4	6.9
Random Effects	23810.0	129.6

5. Concluding Remarks

In this paper we have presented a modulated Poisson process model for describing call center arrival data and developed its Bayesian analysis using Gibbs sampling. In order to incorporate potential heterogeneity in advertisements, we have considered a random effects type extension of the model and discuss Bayesian model comparison.

The proposed approach and the models have been implemented using real call arrival data on 142 advertisements and type of insights that can be obtained from the Bayesian analysis have been illustrated. The analysis of the data has shown that the effectiveness of advertisements deteriorates by time and certain offer types increase the

call volume. The analysis using the random effects model suggests presence of heterogeneity in call volume generating ability of the advertisements that can not be described by the covariates. Our comparison of the fixed and random effects type models using DIC has shown a strong support for the random effects model.

The presented analysis has assumed a power law model for the baseline cumulative intensity. The appropriateness of this choice for a given call volume data can be investigated by considering alternate models for the baseline cumulative intensity and using DIC to compare these models. Alternatively, a semi-parametric modeling strategy can be used by treating the baseline cumulative intensity $\Lambda_0(t)$ as a nonparametric form and by specifying a parametric prior for β. This is an area for future research.

Another area of further study is the formal incorporation of prior opinion into the analysis of modulated Poisson process model. An extension of the methodologies presented in Campodonico and Singpurwalla (1995), for using expert opinion in analysis of point process models, can be considered for this purpose.

References

Allenby, G. M. and Rossi, P. E. (1999). Marketing Models of Consumer Heterogeneity. *Journal of Econometrics*, Vol. 89, pp. 57-78.

Andrews, B. H. and Parsons, H. L. (1993). Establishing Telephone-agent Staffing Levels through Economic Optimization. *Interfaces*, Vol. 23, pp. 15-20.

Andrews, B. H. and Cunningham, S. M. (1995). L. L. bean Improves Call-Center Forecasting. *Interfaces*, Vol. 25, pp. 1-13.

Avramidis, A. N., Deslauriers, A. and L'Ecuyer, P. (2004). Modeling Daily Arrivals to a Telephone Call Center. *Management Science*, Vol. 50, pp. 896-908.

Campodonico, S. and Singpurwalla, N. D. (1995). Inference and Predictions for Poisson Point Processes Incorporating Expert Knowledge. *Journal of the American Statistical Association*, Vol. 90, pp. 220-226.

Cox, D. R. (1972a). The Statistical Analysis of Dependencies in Point Processes. In *Stochastic Point Processes*. Ed. P. A. W. Lewis, pp. 55-66, New York Wiley.

Cox, D. R. (1972b). Regression Models and Life Tables. *Journal of the Royal Statistical Society, Ser. B*, Vol. 34, pp. 187-220.

Dellaportes, P. and Smith, A.F.M. (1993). Bayesian Inference for Generalized Linear and Proportional Hazards Models via Gibbs Sampling. *Applied Statistics*, Vol. 42, pp. 443-459.

Gans, N., Koole, G. and Mandelbaum, A. (2003). Telephone Call Centers: Tutorial, Review and Research Prospects, Manufacturing & Service Operations Management, Vol. 5, pp. 79-141.

Gelfand, A. E. and Smith, A. F. M. (1990). Sampling-Based Approaches to Calculating Marginal Densities, *Journal of the American Statistical Association*, Vol. 85, pp. 398-409.

Gelfand, A. E. (1999). Approaches for Bayesian Semi-parametric Regression. In *Asymptotics, Nonparametrics and Time-Series*. (Eds. M. L. Puri and S. Ghosh), Marcel Dekker.

Gilks, W. and Wild, P. (1992). Adaptive Rejection Sampling for Gibbs Sampling. *Journal of the Royal Statistical Society, Ser. B*, Vol. 41, pp. 337-348.

Jongbloed, G. and Koole, G. (2001). Managing Uncertainty in Call Centers using Poisson Mixtures. *Applied Stochastic Models in Business and Industry*, Vol. 17, pp. 307-318.

Kass, R. E. and Raftery, A. E. (1995). Bayes Factors. *Journal of the American Statistical Association*, Vol. 90, pp.773-777.

Lindley, D. V., and Smith, A. F. M. (1972). Bayes Estimates for the Linear Model (with discussion). *Journal of the Royal Statistical Society, Ser. B*, Vol. 34, pp. 1-41.

Lindley, D. V. (1983). Reconciliation of Probability Distributions. *Operations Research*, Vol. 31, pp. 866-880.

Merrick. J. R., Soyer, R. and Mazzuchi, T. A. (2004). Are Maintenance Practices for Rail Road Tracks Effective ? Under review *Journal of the American Statistical Association.*

Rossi, P. E. and Allenby, G. M. (2003). Bayesian Statistics and Marketing. *Marketing Science*, Vol. 22, pp. 304-328.

Sinha, D. (1993). Semiparametric Bayesian Analysis of Multiple Time data. *Journal of the American Statistical Association*, Vol. 88, pp.979-983.

Spiegelhalter, D., Thomas, A., Best, N., and Gilks, W. (1996). Bayesian Inference Using Gibbs Sampling Manual (version ii). MRC Biostatistics Unit, Cambridge University.

Spiegelhalter, D. J., Best, N. G, Carlin, B. P. and van der Linde, A. (2002). Bayesian measues of model complexity and fit, *Journal of the Royal Statistical Society, Series B*, Vol. 64, pp. 1-34.